

Statement by Raj Reddy
Co-Chair, President's Information Technology Advisory Committee
March 8, 2000
Senate Commerce, Science, and Transportation Committee
Subcommittee on Communications

Introduction

Mr. Chairman and Members of the Subcommittee, thank you for this opportunity to testify about important research and development efforts aimed at increasing Internet security and protecting our Nation's Information Infrastructure.

My name is Raj Reddy, and I am the Herbert A. Simon University Professor of Computer Science and Robotics at Carnegie Mellon University. I also serve as Co-Chair of the President's Information Technology Advisory Committee, commonly known as PITAC. In the PITAC's February 1999 report to the President, "Information Technology Research: Investing in Our Future," we highlighted the need for increased investment in network security, as well as other important research areas. Today, on behalf of PITAC, I will provide you with insight into the state of Internet security in our country and outline some of the PITAC recommendations that will help our Nation build and support a more reliable, available, secure, and scalable Internet. I will also present my personal views on an R&D strategy for developing and demonstrating highly dependable networks.

Background

While advances in information technology have created unprecedented economic growth and transformed our lives in thousands of positive ways, weaknesses still exist that enable malicious hackers to disrupt Internet service and overload popular Web sites. An analysis of the highly visible disruptions to Internet access reveals a wide range of causes, including denial of service attacks from malicious hackers using insecure hosts infected with "zombie" diseases (Yahoo!), software bugs (Ameritrade), insecure configurations (Schwab), change management (E-trade), and security loopholes (Hotmail, Melissa).

PITAC shares Congress' concern about these recent hacker attacks. In our report to the President, we observed that "the Internet is growing well beyond the intent of its original designers and our ability to extend its use has created enormous challenges. As the size, capability, and complexity of the Internet grows, it is imperative that we do the necessary research to learn how to build and use large, complex, highly-reliable, and secure systems... It is therefore important that the Federal government undertake research on topics ranging from network reliability and bandwidth, to robust, reliable, secure ways to deliver and to protect critical information." In our report, we recommended a research agenda to help ensure the survivability of our information infrastructure in the face of malicious attacks or viruses, equipment or software failures, and overload. Before I discuss the specifics of the R&D agenda

for Internet security, I would first like to briefly summarize the findings and recommendations of our report.

The PITAC Report Findings and Recommendations

The PITAC was established pursuant to the High Performance Computing Act of 1991 and was tasked to look at a number of issues in high performance computing and communications. After a detailed review of the Federal IT R&D programs, we concluded that U.S. leadership in IT provides an essential foundation for promoting economic growth, education and research, environmental stewardship, public health, and national security. We also concluded that there has been an erosion of support for long-term fundamental research in IT and that current research is too focused on near-term problems linked to agency missions. Our Committee recommended that the Federal government create a strategic initiative for long-term R&D and increase funding for IT R&D by \$1.4 billion by fiscal year 2004 over the fiscal year 1999 base programs funding level. Our report recommended a balanced research agenda, with priority for the following areas:

- * Software: Methods for efficiently creating and maintaining high-quality software of all kinds and for ensuring the reliability of the complex software systems that now provide the infrastructure for much of our government and our economy.
- * Scalable Information Infrastructure: Techniques for ensuring that the National Information Infrastructure consisting of communications systems, the Internet, large data repositories, and other emerging systems-is reliable and secure, and can grow gracefully to accommodate the massive numbers of new users (perhaps billions) and applications expected over the coming two decades.
- * High End Computing: Continued invention and innovation in the development of fast, powerful computing systems and the accompanying communication systems are needed to implement critical science, engineering, and business applications ranging from aircraft design to weather and climate modeling.
- * Social, Economic, and Workforce Implications of IT: Research directed towards better understanding the sociological and economic impacts of innovations in information technology and toward growing the workforce to meet the national need for information technology professionals.

Our recommendation for research to support a scalable information infrastructure included topics to enable the survivability of our networks and information. Survivability means that services will be available when needed and information will be delivered in a timely fashion. The recommended research agenda includes:

- * Authentication and security mechanisms for a large, heterogeneous, and evolving

infrastructure

- * Mechanisms for detecting system intrusion and information software corruption
- * Mechanisms for detecting, mitigating, responding to, and recovering from, or for preventing, human error in the creation and use of the infrastructure
- * Mechanisms for assuring information quality
- * Scalable information and service replication strategies
- * Mechanisms for monitoring services to ensure correct operation within given quality-of-service bounds
- * Repositories for guaranteed long-term preservation of information

Our report recommendations have received strong bi-partisan support and we were encouraged by the \$235 million increase for IT R&D appropriated in this year's budget. The President's fiscal year 2001 budget proposes an increase of nearly \$600 million in IT R&D in a balanced research program that addresses the recommendations in the PITAC report. Proposed funding includes networking and software research directed towards technologies to enable more secure, reliable, and dependable networks. The PITAC applauds the Senate's past support and leadership for IT R&D and hopes the Senate will support the full set of research priority areas recommended in our report.

The PITAC report provides broad concepts for a balanced IT R&D program. While we recognized the importance of network security, reliability, and dependability, we did not develop a detailed R&D agenda for Internet security. Our recommendations cover a range of important topics to be addressed, rather than proposals for specific research projects.

The Impact of Internet Downtime on Businesses and Society

Denial of service happens when the network fabric is overloaded through intentional and unintentional ("legal") overloading of the system with too many requests. This is analogous to a large number of people calling California in the event of an earthquake report, or a computer calling a phone continuously thereby blocking anyone else getting through in case of an emergency.

The cost of denial of service and overloading can be substantial. The Yankee Group estimates that the online industry may have lost \$1.2 billion in revenue from the Web site attacks earlier this month. (WSJ, Feb 24, 2000). A Gartner Group study showed that the average cost of downtime in brokerage operations is about \$6.5 million per hour! According to the Boston-based market research firm, \$29 million in refunds were paid out by MCI to customers affected by the 10 day outage of its frame relay network in August 1999. Three thousand companies were affected. (Online News, 10/28/99). eBay paid \$3.9 million in credits to its customers for the service outage that halted bidding completely at its popular service for an unprecedented 22 hours in June 1999. Distributed network sites can lose \$20,000 to \$80,000 per hour. (Computer Reseller News, 1998). At a cost of \$80,000 per hour, the average company will lose \$7.1 million per year in centralized network downtime.

These costs are expected to increase as companies incur indirect costs in the form of lawsuits, regulatory scrutiny, impact on brand name and public image, loss of customer base, lower employee morale and productivity, and higher employee stress.

The impact on businesses of system outage can be even more devastating. In an April 1999 survey of consumers, research firm Jupiter Communications found that 46 percent leave a preferred site if they experience technical or performance problems. Statistics from McGladrey and Pullen show that for every five organizations affected by a disaster, two will be unable to maintain their critical business functions and make a recovery. Of the remaining three, one will not survive the next two years. In fact, a company that experiences a computer outage lasting more than 10 days will never fully recover financially ("Disaster Recovery Planning: Managing Risk and Catastrophe in Information Systems" by Jon Toigo).

According to Cahners in-stat group, Internet downtime hits businesses financially, (http://www.instat.com/abstracts/ia/1999/is9906sp_abs.htm), affecting direct revenue/customer base, compensatory payments, inventory cost, and depreciation of capital. It also affects business in ways not seen on the balance sheet, such as market capitalization loss, employee downtime, and delays to market items that may prove more financially damaging than the explicit losses associated with an outage. The report "Data Failure: the financial impact on Internet business" quantifies the real-cost damages for site outages based on SEC filings and publicly released information. The report compares two e-commerce business models and illustrates how much is at stake in the event of data failure.

Steps Towards a Secure and Dependable Internet

Many of the problems of Internet access can be avoided by taking some simple common sense precautions. For example:

Online businesses can:

- * Educate users on cyber hygiene, security tools, and procedures such as use of firewalls, intrusion detection systems, anti-virus software, automatic daily disinfecting tools, etc.
- * Discourage masquerading and spoofing attacks by ensuring that network traffic exiting from the local area network of an organization carries the address consistent with the valid set of addresses for that organization.
- * Protect against inside hacker risk by providing backup and retrieval from an off-site storage service provider. Disaster tolerance backup facilities are offered by many suppliers. Such services guarantee constant availability of data in the face of technical or natural catastrophe, including surge capabilities for unplanned swells in site traffic.
- * Provide 24 hour-per-day, 7 day-a-week physical security to central facilities and server farms. Alternatively, use the backup and retrieval from an off-site storage service as described in the previous bullet.

Industry can:

- * Release hardware and software that prevents insecure configurations and provide tools for intrusion detection.
- * Re-engineer operating systems and applications to make them immune to the effects of viruses and other forms of malicious code.
- * Identify and close the security loopholes and backdoors by working with major vendors to provide access to the source code and encourage open source movement.
- * Develop and deploy a secure communications infrastructure that can be used by network operators and Internet service providers to enable real-time collaboration when dealing with attacks.

Many of the common sense measures listed above depend on the voluntary compliance of more than a 100 million Internet users and organizations that provide Internet service. However, history has shown us that compliance failures will occur, either unintentionally or maliciously. Rather than leaving the Internet vulnerable because a few persons or organizations are careless or reckless, we should develop an information infrastructure that is not dependent on voluntary compliance of security practices and policies.

Personal Views on a Strategy for a National Self Healing Network Testbed

I would now like to make some personal observations and make a specific recommendation for creating a national self healing network testbed. The PITAC recommended an aggressive new program in networking research, including network security. We also recommended expanded research to explore ways that laws protecting privacy, intellectual property, and other rights are extended effectively into this new media. We continue to support increased funding in these critical areas.

The PITAC is currently reviewing federal research plans and will be issuing new recommendations later this year. Since these new recommendations are not available, I would like to present my personal views on logical next steps.

By now we understand the sources of highly publicized Internet crashes: malicious hacker attacks and "legal" users overloading popular web sites. Many of the remedies require straightforward implementation of known solutions, either administrative or legal. However, herein lies the problem – we simply cannot depend on every system to be properly administered or every person to behave as desired. Instead, we should strive to develop an Internet infrastructure in which it does not matter if someone is careless or reckless. In my view, one of the key goals of networking research over the next few years should be development of a "self healing" network. A self healing network would work similar to the human immune system. It would

constantly monitor the system (in this case, the network), analyze what is in the system, and if it finds something wrong within the system, immediately begin actions to remedy the problem. A self healing network would be capable of self-monitoring, self-diagnosing and self-repairing. To accomplish this, we should establish a national network testbed that can be used to develop and demonstrate what I will refer to as an "ultra-dependable Internet." This is similar to an ultra-high speed network, but with the focus on dependability rather than speed.

I will use the phrase "dependable Internet" to specifically include attributes such as reliability, availability, and scalability in addition to security. The operative issue is not "security" as interpreted narrowly in the research circles but rather "how to create a dependable Internet Infrastructure" that is as reliable as the current telephone system. By dependable, I mean a system ("as if my life depended on it") that is:

- * **reliable**, i.e., always up, accessible, accurate, and consistent,
- * **available**, i.e., a system with no world-wide-wait and a response time of under 200 milliseconds most of the time,
- * **scalable**, i.e. an infrastructure capable of scaling to a billion simultaneous users and a trillion inter-connected devices, and
- * **secure**, i.e. no fear of loss of privacy and immunity to sniffing and spoofing.

The goal of a self healing network is to provide mechanisms for detecting unauthorized use of networking equipment, tracking inappropriate uses, and identifying the individuals using networks for malicious intent, without compromising individual rights to privacy and security on the network. Over the years we have found ways to balance privacy and security in traditional commerce. Applying these precedents to the new networked world will require combining the skills of technologists and people knowledgeable of the legal, economic, and social issues. Clearly this is an enormous challenge, but I believe it is a critical national research challenge and deserves and appropriate response.

A Self Healing Network

A self healing network is one which continuously monitors all the traffic within the system (every packet entering the system is validated before it can proceed) with a view to detect and disable abnormal traffic patterns. It is predicated on using "software agents" capable of self-monitoring, self-diagnosis, and self-repair much as the human immune system uses (distributed) anti-bodies to disable antigens and restore balance in the human body. Just as in human systems where a few people may get sick some of the time, but society as a whole continues to function, we may accept an occasional denial of service as long as most users are able to access most of the web sites without any degradation of service.

Self monitoring within the Internet core fabric requires agents capable of continuous and autonomous monitoring of "packet" traffic using "software sensors." "Self repair agents" undertake a set of autonomous corrective actions against the offending source that is generating the unusual traffic by dropping the packets or limiting it to a "fair share" the number of packets entering the fabric. The work of these agents and the humans tracking network security could be helped if the new generation of routers add information packets that make it easier to detect malicious patterns of use and to track the attacks to their source.

The proposed self healing network will add to the packet handling overhead at each router in the fabric and has the potential to make the system slower, waste bandwidth, and compromise privacy. At first blush, this requirement appears to be impractical, as the Internet is expected to handle trillions of packets every day and would require expensive retrofitting of the existing commercial Internet Service Providers (ISPs). However, such a transition is not only essential to the future economic growth and security of the nation, but also practical given the expected exponential advances in processor, memory, and optical networking technologies. The expected additional overhead in packet handling will be ameliorated by better algorithms, exponential improvements in processor (predicted by Moore's law), memory, and bandwidth technologies and increasing locality of Internet traffic patterns ("Internet is global and the traffic is local").

In addition to the research needed to develop terabit networks, faster routers, efficient algorithms, and distributed computation techniques, research will also be needed for data warehousing of meta-data contained in packet headers, data mining of this data to establish statistical parameters that can be used to classify normal and abnormal traffic requests, and repair strategies for generating a signal (analogous to the busy signal used in voice telephony) to sites making abnormal requests without prior arrangement for surge capacity.

Conclusion

In conclusion, creating a dependable Internet infrastructure that is as dependable as telephone service is essential to the future economic growth and security of the nation. It is possible to create a system capable of achieving these goals while ensuring absolute protection of personal privacy and without major reductions in networking speed. Indeed, rapid advances in computing power and networking speed should make the new security systems nearly invisible to users. The main challenge is to find the right balance between having a dependable Internet infrastructure without compromising the ease of use by non-experts and protecting the privacy of the individuals connected to the infrastructure. To accomplish this will require both new research ideas and the uniform application of known and new ideas across the Internet infrastructure. It makes sense to apply the creative energies of academe to these social problems.

Development of networks capable of meeting our goals for security and privacy will only happen with a concerted research investment supported by both government and industry. One strategy would be to support a network testbed designed with the specific goal of evaluating innovative

strategies for network protection – including commercial concepts. Such a testbed would provide useful networking services and at the same time let commercial operators and government research organizations evaluate advanced networking security concepts.

It is estimated that market capitalization of Internet based industries created since 1990 is more than a trillion dollars resulting in capital gains taxes of more than \$200 billion to the nation. Investing a small fraction of this national income in research towards creating a self healing Internet will ensure the continuation of this engine of growth!

Acknowledgements

This paper has benefited from the comments and suggestions from several PITAC members: Jim Gray, Irving Wladawsky-Berger, Vint Cerf and Bob Kahn and from other colleagues: Anish Arora, V.S. Arunachalam, Ed Lazowska, and Rich Pethia. Please send comments to rr@cmu.edu.